DDR2: The Next Generation Main Memory

By Jimmy Ma

Introduction

Today's memory architecture shows significant improvements when compared to the days of Fast Page Mode (FPM) and Extended Data Out (EDO). The industry has shifted gear from an asynchronous world of FPM and EDO to a synchronous one. The majority of memory makers now use Synchronous Dynamic Random Access Memory (SDRAM). What was once only considered as an interim technology, SDRAM has made a significant impact on the computing market.

The first generation of synchronous memory (PC66) uses single-ended technology. And it didn't take long before PC100 and PC133 debuted with an increase in operating frequency ranging from 66MHz to 100MHz and 133MHz for PC100 and PC133, respectively. PC100 and PC133 retained the single-ended technology used in PC66. But as demands for higher speed and lower power consumption became more demanding, single-ended signaling began to reach its limitations. Once again, the memory industries have shifted gear to meet the new demand and a newer technology emerged with promises of lower power consumption and faster operating frequency. It was DDR.

DDR is an acronym for Double Data Rate, a memory technology that has evolved from its predecessor, the single-ended data rate PC100 and PC133. Double Date Rate, as its name suggests, uses technology of acquiring data on both the rising and falling edge of the clock. The technology has been well adopted in the computing and consumer markets such as workstation, desktop PC, video cards, and set-top-boxes. But with the computing and consumer markets continuing to evolve and processor speeds continuing to rise, DDR is migrating to the next generation of memory architecture to support the performance challenge. This next generation of memory is DDR2. Being the next generation of high performance memory solution, DDR2 boasts an initial clock speed of 400MHz with a data transfer rate of 3.2Gbit/s. When compared to PC100, DDR2 initial speeds have quadrupled the data rate.

Table 1 provides an overview of the performance characteristics of past and future of memory technologies. Table 2 provides a more detail overview of past and future of memory technologies:

Table 1. Memory Technology Past and Future

SDR	3.3V	PC100	N/A	100MHz	0.8Gb/s
SDR	3.3V	PC133	N/A	133MHz	1.06Gb/s
DDR	2.5V	PC1600	DDR200	100MHz	1.6Gb/s
DDR	2.5V	PC2100	DDR266	133MHz	2.1Gb/s
DDR	2.5V	PC2700	DDR333	166MHz	2.7Gb/s
DDR*	2.5V	PC3200	DDR400	200MHz	3.2Gb/s
DDR-II	1.8V	PC23200	DDR2-400	200MHz	3.2Gb/s
DDR-II	1.8V	PC24300	DDR2-533	266MHz	4.3Gb/s
DDR-II	1.8V	PC25300	DDR2-667	333MHz	5.3Gb/s

Note: * = Proposed, but not yet approved by JEDEC
      SDR = Single Data Rate
      DDR = Double Data Rate

www.pericom.com

Table 2. Electrical and Physical Comparison for Memory Technology Past and Future

Clock Frequency	100 / 133	100 / 133 / 166 / 200	200 / 266 /333
Data Rate	100 / 133	200 / 266 / 333 / 400	400 / 533 /667
Voltage	3.3V	2.5V	1.8V
Densities	32MB – 512MB	128MB – 2GB	256MB – 4GB
Internal Banks	4	4	4 to 8
Termination for PLL in Registered DIMM	Series Termination	120	120
Type of Signaling	LVTTL	SSTL_2	SSTL_1.8
Data Strobes	Single-Ended	Differential	Single-ended or Differential
Memory Bus Width	64 Bit	64 – Bit	64 - Bit
Signal Type	Single-ended	Differential	Differential
Module Pin Count	100-Pin Unbuffered DIMM 168-Pin Registered DIMM 168-Pin Unbuffered DIMM 144-Pin SODIMM	100-Pin Unbuffered DIMM 184-Pin Registered DIMM 184-Pin Unbuffered DIMM 172-Pin MicroDIMM	240-Pin Registered DIMM 240-Pin Unbuffered DIMM 200-Pin SODIMM
Package	TSOP, FBGA, VFBGA	TSOP, FBGA	FBGA

PERICOM®
Semiconductor Corporation

What's new with DDR2?

Being the next generation of memory solution, the DDR2 architecture implemented multiple features to distinguish itself from DDR1. The first and most obvious feature is higher operating frequency while operating at lower supply voltage. With an initial clock speed of 400MHz, it continued where DDR1 has left off. With that, DDR2 promises to increase clock speeds up to 667MHz which means transfer rates will increase to a whopping 5.3Gb/s. This is an enormous leap from PC100 with a transfer rate of only 0.8Gb/s. Amazingly the core frequency still remains at 100MHz. The I/O buffer in PC100/133 and DDR1 are synchronous with the core with the exception that the protocol varies. In the case of DDR2, the I/O buffers are clocked twice the rate of the core frequency. What this means is DDR1 prefetches 2 bits while DDR2 is required to prefetch four bits. Figure 1 shows a graphical view of the core and I/O frequency for SDR, DDR1, and DDR2.
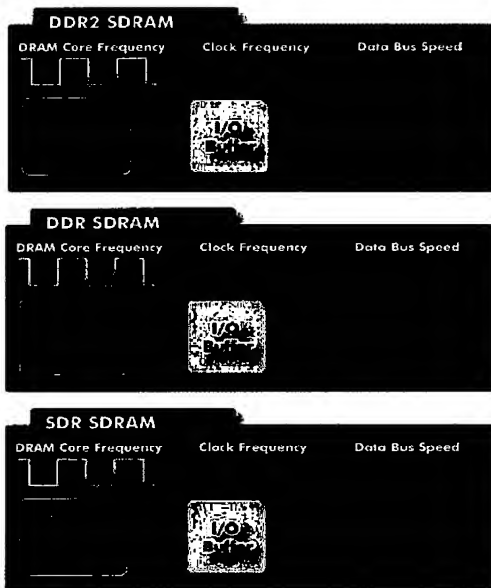

Figure 1. Memory core frequency

www.pericom.com

PERICOM®
Semiconductor Corporation

A Closer Look at Registered DIMM

Similar to its predecessor, DDR2 comes in both unbuffered and registered (buffered) types of DIMM (Dual Inline Memory Module). Unbuffered DIMMs are mainly used in PC desktops and notebooks, whereas, registered DIMM are mainstream in high-end platforms such as severs, datacom and telecom systems. But as system design becomes more demanding, registered DIMM may one day emerge onto the consumer desktop and notebook PCs. A complete block diagram of how a registered DIMM operates in a server is shown in Figure 3.
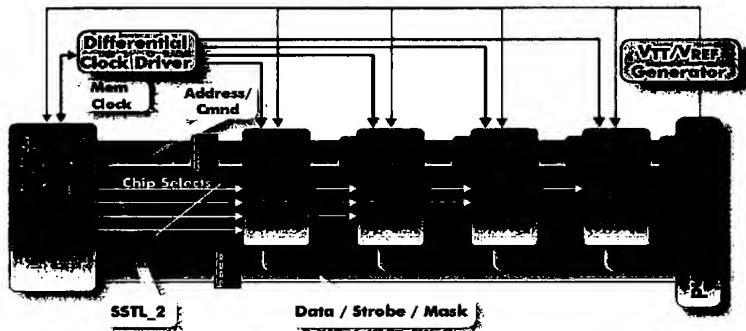

Figure 3. Overall block for a Registered DIMM

An unbuffered DIMM and a registered DIMM are somewhat similar. However, a closer look at a registered DIMM shows that there are typically three additional components: a Phase-Lock Loop (PLL) zero-delay clock buffer and two registers for a double-sided DIMM, or one register for a single-sided DIMM. The mainstream DIMM modules are double-sided DIMMs. Double-sided DIMMs have memory chips populated on both sides of the module. As mentioned earlier, registered DIMMs have a PLL. The PLL is necessary in order to distribute additional clock signals while removing any delay generated by the device itself and the traces, hence the name zero-delay buffer. A more technical explanation on how PLL generates a zero delay is provided below. The other two components, the registers, are used for address lines and control signals. The on-board registers on buffered DIMMs help reduce the system's loading caused by the DIMMs on the server board. As a result, this will allow the system to load larger capacity DIMMs and increase the number of DIMMs to be populated onto the server board. In other words, registered DIMM offers reduced system loading by allowing the MCH (Memory Controller Hub) to see it as a single load. Figure 4 is focused on the PLL and registers connection.
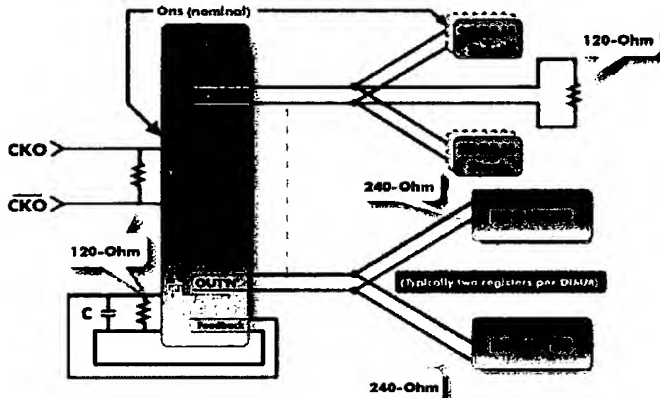

**Figure 4. Architecture structure of a registered DIMM**

Tuning for Perfect Time

As frequencies increase, the period of each cycle becomes shorter leaving little room for any error. However, with the addition of a PLL, it provides the ability to make adjustments on the stringent timing requirement to increase the system-timing margin. The main characteristic of a PLL clock buffer compared to its counterpart, a non-PLL clock buffer, is the ability to adjust the system clock timing by altering the delay in the feedback loop.

The concept of a zero-delay buffer allows the timing between the input clock, CLK_IN signals, and the output of the clock at the destination of the memory chip, MEM_IN signals, to be adjusted to have the most optimal timing. This type of tuning will allow the PLL to generate a leading skew or a lagging skew to compensate for any delay generated by the device itself and the traces. By leading skew, it indicates that the MEM_IN comes earlier than CLK_IN. Whereas by lagging skew, it indicates that the MEM_IN will come later than CLK_IN. The mechanism of generating a leading or lagging skew becomes very valuable to memory module makers since it allow the manufactures to make adjustments to the timing in order to meet any system timing requirements.

Feedback Circuit

The timing adjustment is done through the feedback loop of the PLL. The feedback circuit of the PLL is the trace and the circuit between the FBIN and FBOUT of the PLL as shown Figure 5. It is recommended that 70% to 85% of the feedback delay should be generated by the trace in the feedback loop. Note that generating more delay will cause the MEM_IN signal to lead CLK_IN. In most cases, due the constraint of board space, only one feedback loop is possible. But with components constantly reducing in size and depending on the board design, it is possible to make a secondary feedback loop. The secondary feedback loop is to provide more accessibility for adjustments once the modules have been fabricated. In situations where the module has been fabricated and the first loop does not provide any suitable timing for the system, it is still feasible to make further adjustments by disengaging the primary trace loop and connecting the secondary trace. The secondary loop can be connected through a zero-ohm resistor as shown in Figure 6. While the feedback trace will generate about 70% to 85% of the delay, the smaller delay is generated through the capacitor. The capacitor should generate a timing delay of no more than 15% to 30% since its purpose is used only to fine tune the timing. The default capacitor value is typically 0pF, but can be increased to 10pF for timing adjustments. It is strongly recommended that the capacitor value should not be too large since it can result in the feedback signal to have smaller swing range and the reduction in the slew rate leading to additional jitter. If the feedback signal has a significant amount of jitter, it will reflect to the output signals thus resulting in the degrading of the DIMMs performance. One important note to take into consideration is that it is always easier to generate delay to have MEM_IN leading CLK_IN once the board has been fabricated. However, if the system requires MEM_IN to lag CLK_IN, it is much more difficult to reduce the delay than to increase the delay since it would require the trace length to be reduced resulting in a re-spin of the board.
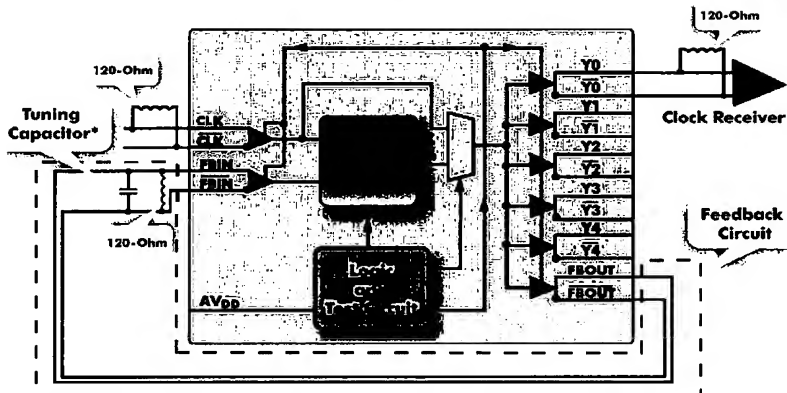
*Note: Tuning capacitor will vary from one system to another

Figure 5. PLL architecture and Feedback circuit
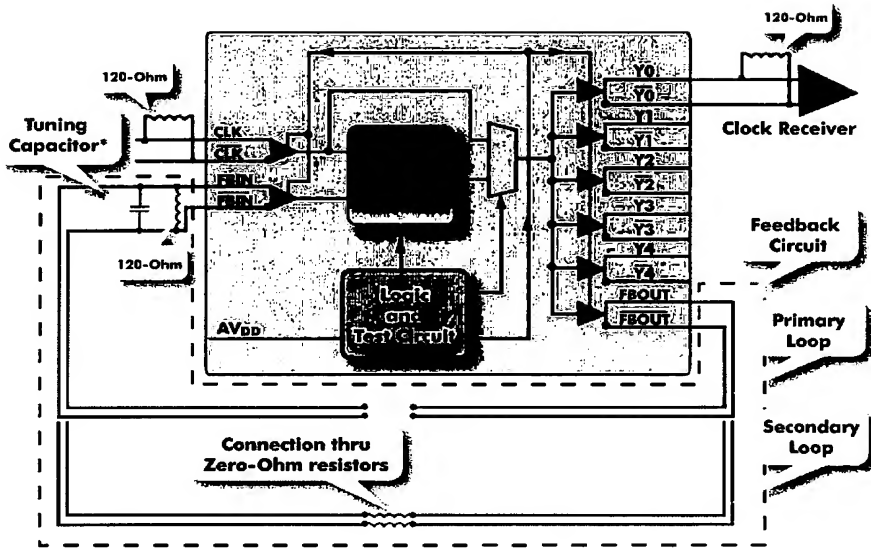
*Note: Tuning capacitor will vary from one system to another

Figure 6. Optional design with PLL secondary feedback loop

An actual application is shown graphically in Figure 7 and 8. In Figure 7, the PLL is set to have a default skew of 0ps between the input clock CLK_IN and the output clock at the memory chip MEM_IN, for simplicity. As we decreased the delay in the feedback loop by reducing the capacitor value, the skew between CLK_IN and MEM_IN becomes a lagging skew (MEM_IN comes later then CLK_IN); Figure 8a. Similarly, to generate a leading skew, MEM_IN ahead of CLK_IN, we increased the capacitor value; Figure 8b. Note that the changes are in a few hundred picoseconds. As mentioned earlier, the capacitor is mainly used to fine-tune the timing between CLK_IN and MEM_IN.

Table 3. Timing adjustment methodologies

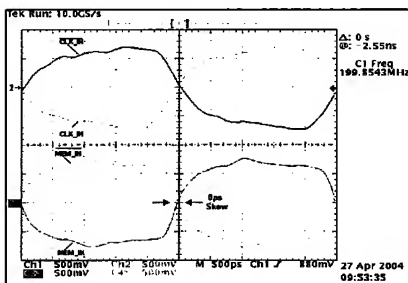|		
MEM_IN Lagging CLK_IN	Decrease the delay in the feedback loop 1. Decrease the feedback capacitor in the FB loop 2. Decrease the FB loop trace length	Figure 8a
MEM_IN Leading CLK_IN	Increase the delay in the feedback loop 1. Increase the feedback capacitor in the FB loop 2. Increase the FB loop trace length	Figure 8b

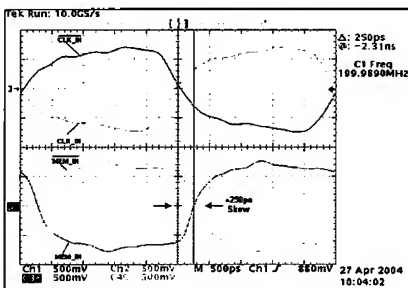**Figure 7. Skew tuned to have 0ps as reference**

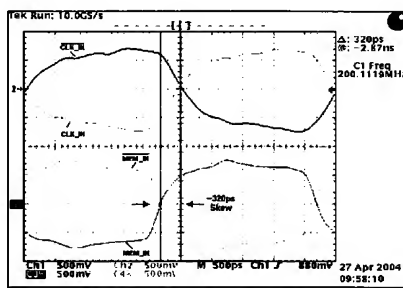**Figure 8a. Reducing the FB capacitor value will generate a lagging skew**

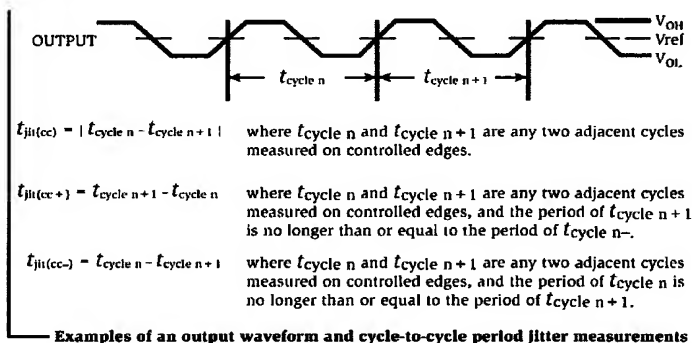**Figure 8b. Increasing the FB capacitor value will generate a leading skew**
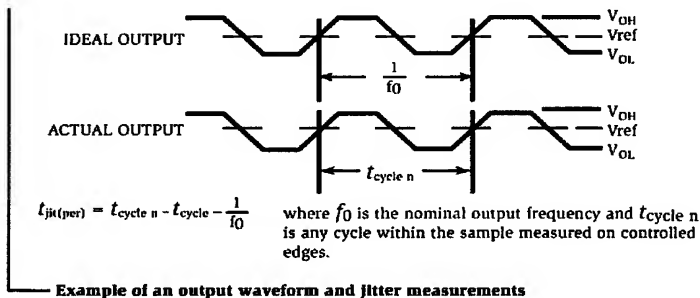
PERICOM®
Semiconductor Corporation

Jitter: What types are there, and how do they affect the system timing?

Since timing is extremely stringent in DDR2, jitter becomes a great concern. The PLL itself will generate some form of jitter. This is just the nature of a PLL. Even the most minute amount of jitter will eat up the system timing and can lead systematic problems. The question to address for system designers might be *what are the major types of jitter to be concerned with and how much the system can tolerate it.* Some of the major jitter parameters to take into consideration when designing with registered DIMM using PLL are addressed below:
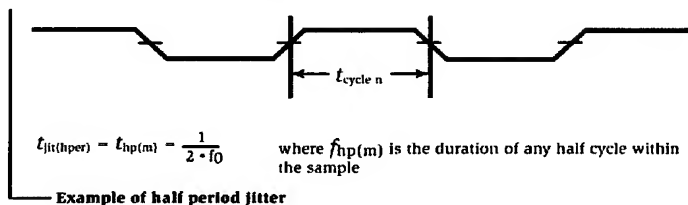
Cycle-to-cycle jitter: The variation in cycle time of a signal between adjacent cycles, over a random sample of adjacent cycle pairs.

$t_{jit(cc)} = |t_{cycle\ n} - t_{cycle\ n+1}|$ where $t_{cycle\ n}$ and $t_{cycle\ n+1}$ are any two adjacent cycles measured on controlled edges.

$t_{jit(cc+)} = t_{cycle\ n+1} - t_{cycle\ n}$ where $t_{cycle\ n}$ and $t_{cycle\ n+1}$ are any two adjacent cycles measured on controlled edges, and the period of $t_{cycle\ n+1}$ is no longer than or equal to the period of $t_{cycle\ n-}$.

$t_{jit(cc-)} = t_{cycle\ n} - t_{cycle\ n+1}$ where $t_{cycle\ n}$ and $t_{cycle\ n+1}$ are any two adjacent cycles measured on controlled edges, and the period of $t_{cycle\ n}$ is no longer than or equal to the period of $t_{cycle\ n+1}$.

Examples of an output waveform and cycle-to-cycle period jitter measurements

Period Jitter: The deviation in cycle time of a signal with respect to the ideal period over a random sample of cycles.

$t_{jit(per)} = t_{cycle\ n} - t_{cycle} - \frac{1}{f0}$ where $f0$ is the nominal output frequency and $t_{cycle\ n}$ is any cycle within the sample measured on controlled edges.

Example of an output waveform and jitter measurements
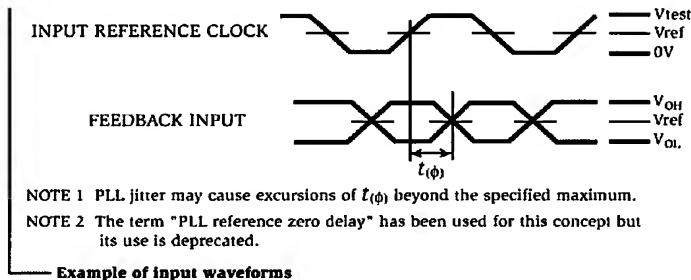
www.pericom.com

Half Period Jitter: The magnitude of the deviation in time duration between half cycle threshold crossing of a signal over a random sample of half cycles.

$$t_{jit(hper)} - t_{hp(m)} - \frac{1}{2 \cdot f_0}$$

where $f_{hp(m)}$ is the duration of any half cycle within the sample

— **Example of half period jitter**

Peak-to-peak Jitter: The sum of cycle-to-cycle jitter, half-period jitter, and period jitter

Dynamic Phase Offset: The incremental phase offset between the input reference clock and the feedback input signal of a PLL resulting from modulations of the input reference clock.

Static Phase Offset: The time interval between similar points on the waveform of the averaged input reference clock and the averaged feedback input signal when the PLL is locked and the input reference frequency is stabled.

INPUT REFERENCE CLOCK — Vtest / Vref / 0V

FEEDBACK INPUT — V_{OH} / Vref / V_{OL}

$t_{(\phi)}$

NOTE 1 PLL jitter may cause excursions of $t_{(\phi)}$ beyond the specified maximum.

NOTE 2 The term "PLL reference zero delay" has been used for this concept but its use is deprecated.

— **Example of input waveforms**

www.pericom.com

To maintain the system-timing requirement, the JEDEC committee has put stringent requirements on the jitter spec.

Table 4: Jitter Specification defined by JEDEC

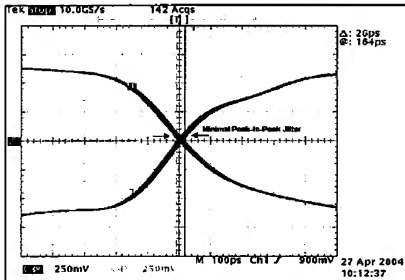|	min		min		min		
Cycle-to-cycle Jitter	-75	+75	-40	+40	-40	+40	ps
Period Jitter	-75	+75	-40	+40	-40	+40	ps
Half-Period Jitter	-100	+100	-75	-75	-75	-75	ps

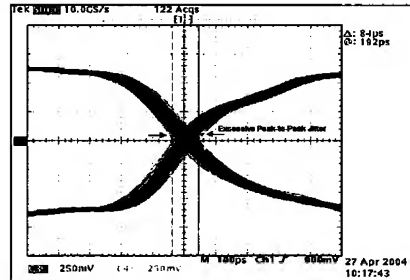**Figure 9a. PLL clock buffer with low peak-to-peak jitter**

Figure 9b. PLL clock buffer with high peak-to-peak jitter

Conclusion: Market Trend

DDR1 has gained credibility for its key innovative performance while being cost effective in all market sectors. DDR2, is the next generation of memory and will continue the journey where DDR1 has left off. The market trend for DDR technology as shown in Figure 10 shows no sign of dying out. Intel, who was a supporter of Rambus, is also hopping onto the DDR bandwagon as is everyone else with chipsets such as the Grantsdale, Alderwoods, and Aliviso to support both DDR1 and DDR2. As the trend in DDR technology continues to grow, Pericom Semiconductor will also grow its product line to support the memory technology. See Table 4.
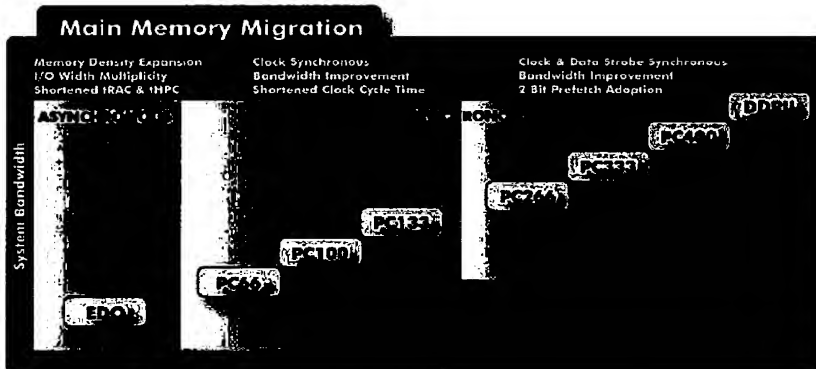

Figure 10. Architecture structure of a registered DIMM

DDR 1

PC1600 / PC2100			
DIMM Configuration		**Planar 9 to 18 loads**	**Stacked 36 Loads**
Standard 1.7" DIMM	PLL	PI6CV857L	PI6CV857
	Register	PI74SSTV16857	PI74SSTV16859
1U Low Profile <1.2" DIMM	PLL	PI6CV857L	PI6CV857
	Register	PI74SSTV16857	PI74SSTV16859 PI74SSTV32852

PC2700			
DIMM Configuration		**Planar 9 to 18 loads**	**Stacked 36 Loads**
Standard 1.7" DIMM	PLL	PI6CV857B	PI6CV857B
	Register	PI74SSTVF16857	PI74SSTVF16859
1U Low Profile <1.2" DIMM	PLL	PI6CV857B	PI6CV857B
	Register	PI74SSTVF16857	PI74SSTVF16859 PI74SSTVF32852

PC3200			
DIMM Configuration		**Planar 9 to 18 loads**	**Stacked 36 Loads**
1U Low Profile <1.2" DIMM	PLL	PI6CVF857	---
	Register	PI74SSTVF16857A PI74SSTVF16859 PI74SSTVF16859A PI74SSTVF32852 PI74SSTVF32852A	---

DDR 2

PC2-3200 / PC2-4300			
DIMM Configuration		**Planar 9 to 18 loads**	**Stacked 36 Loads**
1U Low Profile <1.2" DIMM	PLL	PI6CU877	---
	Register	PI74SSTU32864 PI74SSTU32864A PI74SSTU32866	---

SDR

PC100 / PC133		
DIMM Configuration	**PC100 / PC133**	
	Planar	**Stacked**
Frequency	100MHz / 133MHz	100MHz / 133MHz
PLL Clock	PI6C2510-133E	PI6C2510-133E
Registered Buffer	PI74ALVC162834F	PI74AVC+16334
Registered Buffer	PI74ALVC162835F	PI74AVC+16836
Registered Buffer	-	PI74AVC16834
Registered Buffer	-	PI74AVC16835

References

JEDEC Standard
PC2-3200/PC24300 DDR Registered DIMM Reference Design Specification
Revision 1.0, Updated: January 30th 2003

JEDEC Standard
PC2700 SDRAM Registered DIMM
Design Specification
Revision 1.0, January 2002

JEDEC Standard
SDRAM Registered DIMM
Design Specification
Revision 1.4, August 2001

JEDEC Standard
Definition of Skew Specification for Standard Logic Devices
JESD65-B
Revision of JESD65-A
September 2003

Micron Technology Inc.
Designline Vol. 12, issue 2
 DDR2 Offers New Feature and Functionality
Jeff Janzen © 2003
Elpida Memory Inc.

www.Elpida.com
Document No. E029E30 (Ver3.0)
Published September 2003, Japan